Supervisory Control of Discrete-event Systems under Attacks

Masashi Wakaiki, Paulo Tabuada, and João P. Hespanha

arXiv:1701.00881v1 [cs.SY] 4 Jan 2017

Abstract

We consider a supervisory control problem for discrete-event systems, in which an attacker corrupts the symbols that are observed by the supervisor. We show that existence of a supervisor enforcing a specification language, in the presence of attacks, is completely characterized by controllability (in the usual sense) and observability of the specification (in a new appropriately defined sense). The new notion of observability takes into account the attacker's ability to alter the symbols received by the attacker. For attacks that correspond to arbitrary insertions/removals of symbols, the new notion of observability can be tested by checking the usual notion of observability for a set of discrete-event systems with appropriately redefined observation maps. Focusing on attacks that replace and/or remove symbols from the output strings, we construct observers that are robust against attacks and lead to an automaton representation of the supervisor. We also develop a test for observability under such replacement-removal attacks by using the so-called product automata.

Index Terms

Discrete-event systems, supervisory control, security

M. Wakaiki is with the Department of Electrical and Electronic Engineering, Chiba University, Chiba, 263-8522, Japan (e-mail: wakaiki@chiba-u.jp).

P. Tabuada is with the Department of Electrical Engineering, University of California, Los Angeles, CA 90095, USA (e-mail: tabuada@ee.ucla.edu).

J. P. Hespanha is with the Center for Control, Dynamical-systems and Computation (CCDC), University of California, Santa Barbara, CA 93106, USA (e-mail: hespanha@ece.ucsb.edu).

This material is based upon work supported by the NSF under Grant No. CNS-1329650. The first author acknowledges Murata Overseas Scholarship Foundation and The Telecommunications Advancement Foundation for their support of this work. The work of the second author was partially supported by the NSF award 1136174.

I. INTRODUCTION

Recent developments in computer and network technology make cyber-physical systems prevalent in modern societies. The integration between cyber and physical components introduces serious risks of cyber attacks to physical processes. For example, it has been recently reported that attackers can adversarially control cars [1] and UAVs [2]. Moreover, the Moroochy water breach in March 2000 [3] and the StuxNet virus attack in June 2010 [4] highlight potential threats to infrastructure systems. An annual report [5] published in 2014 by the German government stated that an attacker tampered with the controls of a blast furnace in a German steel factory.

We study supervisory control for Discrete-event systems (DESs) under adversarial attacks. DESs are dynamic systems equipped with a discrete state space and an event-driven transition structure. Such models are widely used to describe cyber-physical systems such as chemical batch plants [6], power grids [7], and manufacturing systems [8]. The objective of this paper is to answer the question: *How do we control DESs if an attacker can manipulate the information provided by sensing and communication devices?* This question encourages system designers to reconsider the supervisory control problem from the viewpoint of security.

Fig. 1 shows the closed-loop system we consider, in which an observation string generated by the plant is substituted by a string corrupted by an adversarial attack. The attack changes the original string by inserting, removing, and replacing symbols, and is allowed to non-deterministically change the same original string to distinct strings. Furthermore, since we cannot foretell what an attacker will do as we design supervisors, we consider a set of possible attacks. The problem we study is how to determine if there exists a supervisor that can enforce the specification notwithstanding the attacks. Whenever such supervisor exists, we also study the problem of how to construct it. This work is inspired by the research on state estimation under sensor attacks for linear time-invariant systems developed in [9]–[11].

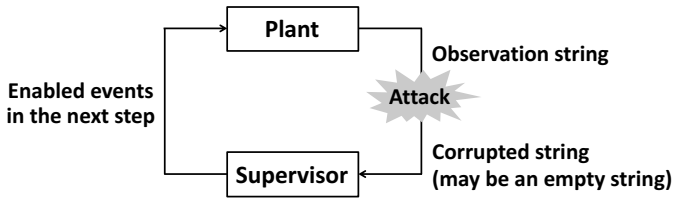

Fig. 1: Closed-loop system under attacks.

Related works: Supervisory control theory has developed frameworks to handle plant uncertainties and faults; see, e.g., [12]–[14] for robust control and [15]–[17] for fault tolerant control. These studies can be used as countermeasures against attacks, but there are conceptual differences between uncertainties/faults and attacks. In fact, uncertainties/faults do not coordinate with harmful intent, whereas attackers can choose their action in order to achieve malicious purposes. For example, the stealthy deception attacks in [18], [19] aim to inject false information without being detected by the controller. Therefore we present a new framework for supervisory control under adversarial attacks.

Several aspects of security have also been explored in the DES literature. One particular line of research aims at studying the *opacity* of DESs, whose goal is to keep a system's secret behavior uncertain to outsiders; see, e.g., [20]–[23] and reference therein. Intrusion detection in the DES framework has been investigated in [24]–[27]. These security methods guarantee the confidentiality and integrity of DESs, but relatively little work has been done towards studying supervisory control robustness against attacks.

An attacked plant can be modeled by a single DES with nondeterministic observations. Supervisory control for such a DES has been studied in [28], [29]. The major difference from the problem setting in these previous works is that in our work, the attacked output, i.e., the nondeterministic observation function, is uncertain. In other words, if G_{A_i} denotes the DES obtained by modeling the plant under an attack A_i, then the problem we consider can be regarded as robust supervision for an uncertain DES in a set of possible models $\{G_{A_1}, \ldots, G_{A_n}\}$, each representing a potential type of attack.

Contributions and organization of this paper: In Section II, after defining attacks formally, we introduce a new notion of observability under attacks, which is a natural extension of conventional observability introduced in [30], [31]. We show that there exists a partial observation supervisor that achieves a given language under attack if and only if the language is controllable in the usual sense and observable according to the new notion of observability introduced in this paper. Moreover, the desired supervisor is obtained explicitly. For attacks that correspond to arbitrary insertions/removals of symbols, the new notion of observability can be reduced to the conventional observability of a set of DESs with appropriately redefined observation maps, and the number of elements in this DES set is the square of the number of possible attacks.

In Section III, we construct an automaton representation of the supervisor derived in Section II. The results in this section are specific to attacks that replace and/or remove specific symbols from

the output string. First, we provide the mathematical formulation of replacement-removal attacks. Then we construct an observer automaton that is resilient against such an attack, extending the result in [30]. Finally, using the observer for each possible attack, we obtain an automaton representation of the desired supervisor.

Section IV is devoted to developing a test for observability under attacks. Attacks are restricted to the replacement and/or removal of symbols in this section as well. Constructing a product automaton in [32], we show how to test observability under replacement-removal attacks in a computationally efficient way. The computational complexity of this observability test is (Number of possible attacks)^2 $\times$ (Complexity of the conventional observability test without attacks), which is the same as in the *insertion-removal* attack case of Section II.

Notation and definitions

The following notation and definitions are standard in the DES literature (see, e.g., [33], [34]).

For a finite set Σ of event labels, we denote by $|\Sigma|$ the number of elements in Σ, and by Σ^* the set of all finite strings of elements of Σ, including the empty string ϵ. For a language $L \subset \Sigma^*$, the *prefix closure* of L is the language

$$\bar{L} := \big\{u \in \Sigma^* : \exists v \in \Sigma^*,\ uv \in L\big\},$$

where uv denotes the concatenation of two strings in Σ^*, and L is said to be *prefix closed* if $L = \bar{L}$. We define the concatenation of languages $L_1, L_2 \subset \Sigma^*$ by

$$L_1 L_2 := \{w_1 w_2 \in \Sigma^* : w_1 \in L_1,\ w_2 \in L_2\}.$$

Let the set of events Σ be partitioned in two sets as $\Sigma = \Sigma_c \cup \Sigma_u$ with $\Sigma_c \cap \Sigma_u = \emptyset$, where Σ_c is called the set of controlled events and Σ_u the set of uncontrolled events. For a language L defined on Σ, a prefix-closed set $K \subset L$ is said to be *controllable* if $K\Sigma_u \cap L \subset K$.

Consider an *observation map* $P : \Sigma \to (\Delta \cup \{\epsilon\})$ that maps a set of events Σ into a set of *observation symbols* Δ (augmented by the empty event ϵ). This observation map P can be extended to map strings of events in Σ^* to strings of observation symbols in Δ^*, using the rules $P(\epsilon) = \epsilon$ and

$$P(w\sigma) = P(w)P(\sigma), \quad \forall w \in \Sigma^*,\ \sigma \in \Sigma.$$

A prefix-closed language $K \subset L$ is *P-observable with respect to L* if

$$\ker P \subset \mathrm{act}_{K \subset L}$$

where $\ker P$ denotes the equivalence relation on Σ^* defined by

$$\ker P := \big\{(w, w') \in \Sigma^* \times \Sigma^* : P(w) = P(w')\big\},$$

and $\mathrm{act}_{K\subset L}$ is a binary relation on Σ^* defined as follows: The pair $(w, w') \in \mathrm{act}_{K\subset L}$ if (and only if) $w, w' \in K$ implies that there does not exist $\sigma \in \Sigma$ such that neither

$$[w\sigma \in K, w'\sigma \in L \setminus K] \text{ nor } [w\sigma \in L \setminus K, w'\sigma \in K].$$

We will omit the underling language L when it is clear from the context.

Consider an automaton $G = (X, \Sigma, \xi, x_0)$, where X is the set of states, Σ is the nonempty event set, $\xi : X \times \Sigma \to X$ is the transition mapping (a partial function), and $x_0 \in X$ is the initial state. We write $\xi(x, \sigma)!$ to mean that $\xi(x, \sigma)$ is defined. The transition function ξ can be extended to a function $X \times \Sigma^* \to X$ according to the following rule:

- For all $x \in X$, $\xi(x, \epsilon) := x$
- For all $x \in X$, $w \in \Sigma^*$, and $\sigma \in \Sigma$,

$$\xi(x, w\sigma) := \begin{cases} \xi(\xi(q, w), \sigma) & \text{if } \xi(q, w)! \text{ and } \xi(\xi(q, w), \sigma)! \\ \text{undefined} & \text{otherwise.} \end{cases}$$

The language generated by G is given by

$$L(G) := \{w \in \Sigma^* : f(x_0, w)!\}.$$

II. Supervised Discrete-event Systems under Attacks

In Section II A, we first introduce attacks on observation symbols and a new notion of observability under attacks. In Section II B, we present the main result of this section, which shows that there exists a supervisor achieving a given language in the presence of attacks if and only if the language is controllable in the usual sense and observable under attacks. Next, in Section II C, we focus our attention on attacks that insert and remove symbols, and show that the new notion of observability under such attacks can be reduced to the conventional observability notion of a set of DESs.

A. Observability under attacks

By an attack, we mean the substitution of an observation string $w \in \Delta^*$ generated by the plant by a corrupted string $y \in \Delta^*$ that is exposed to the supervisor (see Fig. 1). The corrupted string

y may differ from the original string w by the insertion, removal, or replacement of symbols. The simplest form of attack could be modeled by a function $y = A(w)$ that maps Δ^* to Δ^*. However, we are interested in more general forms of attacks where the attacker is allowed to non-deterministically map the same original string $w \in \Delta^*$ to distinct strings $y \in \Delta^*$, in order to make the task of the supervisor more difficult. We thus model attacks by a set-valued function $A : \Delta^* \to 2^{\Delta^*}$ that maps each original string $w \in \Delta^*$ to the set $A(w) \subset \Delta^*$ of all possible corrupted strings y. Note that the supervisor receives one of the strings in the set $A(w)$, not $A(w)$ itself. The attack map $A_{\text{id}} : \Delta^* \to 2^{\Delta^*}$ that assigns to each string $w \in \Delta^*$ the set $\{w\}$ containing only the original string w can be viewed as the absence of an attack.

When we design supervisors, attack maps may be uncertain due to lack of knowledge of which sensors are attacked. In this paper, we therefore consider a set of possible attacks $\mathcal{A} = \{A_1, \ldots, A_n\}$, and we are interested in the following scenario: We know the attack set $\mathcal{A}$ in advance, and that only one attack in the set $\mathcal{A}$ is conducted. In other words, the attacker is not allowed to switch between the attacks in the set $\mathcal{A}$. However, when we construct a supervisor, we do not know which attack actually occurs, and hence the aim is to design a robust supervisor with respect to all attacks in $\mathcal{A}$.

Example 1: Consider the language $L(G)$ generated by the automaton G shown in Fig. 2a. We investigate the observability under attacks of the specification language K generated by the automaton G_K in Fig. 2b. The difference between G and G_K is an event c from x_1 to x_3. The purpose of supervisory control here would be to avoid this "shortcut". We consider the observation map P defined by

$$P(\sigma) = \begin{cases} \sigma & \text{if } \sigma \in \Delta := \{a, b, d\} \\ \epsilon & \text{otherwise,} \end{cases}$$

and three attacks A_1, A_2, A_3 defined by $A_1(\sigma) = \Delta$ for all $\sigma \in \Delta$,

$$A_2(a) = \{a, b\}, \qquad A_2(b) = \{b\}, \qquad A_2(d) = \{a, d\}$$

$$A_3(a) = \{a\}, \qquad A_3(b) = \{b\}, \qquad A_3(d) = \{\epsilon\}.$$

The attack A_1 replaces each output symbol arbitrarily, so the supervisor knows from the output only whether an observable event occurs. The attack A_2 can replace the symbol a by b and the symbol d by a, respectively, whereas the attack A_3 always erases the symbol d. The goal is then to design a supervisor that enforces the specification G_K, without knowing which of the three attacks is taking place. We shall return to this example later in the paper. □

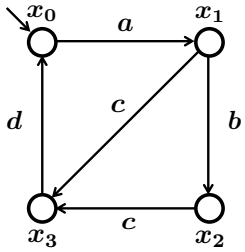

(a) System automaton G.

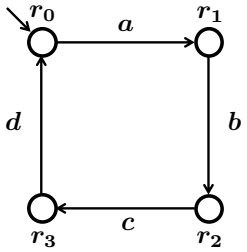

(b) Specification automaton G_K.

Fig. 2: Automata in Example 1.

For simplicity of notation, we denote by $AP : \Sigma^* \to 2^{\Delta^*}$ the *attacked observation map* obtained from the composition $AP := A \circ P$. We introduce a new notion of observability under a set of attacks, which can be seen as a direct extension of the conventional observability notion introduced in [30], [31].

Definition 2 (Observability under attacks): Given an attack set $\mathcal{A}$, we say that a prefix-closed language $K \subset L$ is P-observable under the attack set $\mathcal{A}$ if

$$R_{A,A'} \subset \mathrm{act}_{K \subset L}, \qquad \forall A, A' \in \mathcal{A}, \tag{1}$$

where the relation $R_{A,A'}$ contains all pairs of strings that may result in attacked observation maps AP and $A'P$ with a common string of output symbols, i.e.,

$$R_{A,A'} := \big\{(w, w') \in \Sigma^* \times \Sigma^* : AP(w) \cap A'P(w') \neq \emptyset\big\}. \tag{2}$$

In view of the definition (2), the P-observability condition (1) can be restated as requiring that, for every $w, w' \in K$,

$$\begin{aligned} &\exists A, A' \in \mathcal{A} \ \text{ s.t. } \ AP(w) \cap A'P(w') \neq \emptyset \\ &\qquad \Rightarrow \ \nexists \sigma \in \Sigma \ \text{ s.t. } [w\sigma \in K, \ w'\sigma \in L \setminus K] \ \text{ or } \ [w\sigma \in L \setminus K, \ w'\sigma \in K], \end{aligned} \tag{3}$$

or equivalently, for every $w, w' \in K$,

$$\begin{aligned} &\exists A, A' \ \text{ s.t. } \ AP(w) \cap A'P(w') \neq \emptyset \\ &\qquad \Rightarrow \ \forall \sigma \in \Sigma, \ w\sigma \notin L \text{ or } w'\sigma \notin L \text{ or } [w\sigma, w'\sigma \in K] \text{ or } [w\sigma, w'\sigma \in L \setminus K]. \end{aligned} \tag{4}$$

In words, observability means that we cannot find two attacks $A, A' \in \mathcal{A}$ that would result in the same observation for two strings $w, w' \in K$ such that $(w, w') \notin \mathrm{act}_{K \subset L}$, i.e., two strings

$w, w' \in K$ such that one will transition to an element of K and the other to an element outside K, by the concatenation of the same symbol $\sigma \in \Sigma$.

First we obtain a condition equivalent to observability under attacks for controllable languages. In the non-attacked case, this condition is used as the definition of conventional observability in the book [34, Sec. 3.7].

Proposition 3: Suppose that the prefix-closed language $K \subset L$ is controllable. Then K is P-observable under the set of attacks $\mathcal{A}$ if and only if for every $w, w' \in K$, $\sigma \in \Sigma_c$, and $A, A' \in \mathcal{A}$, the following statement holds:

$$[AP(w) \cap A'P(w') \neq \emptyset,\ w\sigma \in K,\ w'\sigma \in L] \ \Rightarrow\ w'\sigma \in K. \tag{5}$$

Proof: We use the necessary and sufficient condition (3) for the specification language K to be observable under attacks.

($\Rightarrow$) Suppose that K is P-observable under $\mathcal{A}$, and suppose that $w, w' \in K$, $\sigma \in \Sigma_c$, and $A, A' \in \mathcal{A}$ satisfy

$$AP(w) \cap A'P(w') \neq \emptyset, \quad w\sigma \in K, \quad w'\sigma \in L.$$

Then (3) directly leads to $w'\sigma \in K$.

($\Leftarrow$) Suppose that (5) holds for all $w, w' \in K$, $\sigma \in \Sigma_c$, and $A, A' \in \mathcal{A}$. From (3), it is enough to show that if $w, w' \in K$ satisfy $AP(w) \cap A'P(w') \neq \emptyset$ for some $A, A' \in \mathcal{A}$, then there does not exist $\sigma \in \Sigma$ such that

$$[w\sigma \in K,\ w'\sigma \in L \setminus K] \ \text{ or } \ [w\sigma \in L \setminus K,\ w'\sigma \in K].$$

Since K is controllable, if $\sigma \in \Sigma_u$ and $w\sigma, w'\sigma \in L$, then $w\sigma, w'\sigma \in K$ from (5). Moreover, for all $\sigma \in \Sigma_c$, if $w\sigma \in K$ and $w'\sigma \in L$, then $w'\sigma \in K$. Exchanging w and w', we also have if $w'\sigma \in K$ and $w\sigma \in L$, then $w\sigma \in K$. This completes the proof. ■

Example 1 (cont.): Consider the language $L(G)$, the specification language $K = L(G_K)$, and the attack sets A_1, A_2, A_3 in Example 1. Let the controllable event set be $\Sigma_c = \{a, c\}$, and the uncontrollable event be $\Sigma_u = \{b, d\}$. It is straightforward to show that K is controllable, and we can use Proposition 3 to verify that K is observable under the attack set $\mathcal{A} = \{A_1\}$. Additionally, K is observable under the attack sets $\mathcal{A} = \{A_2\}$ and $\mathcal{A} = \{A_3\}$, but is not observable under $\mathcal{A} = \{A_2, A_3\}$. In fact, if we define $w := abcda$ and $w' := wb$, then

$$A_2P(w) = \{abaa, abab, abda, abdb, bbaa, bbab, bbda, bbdb\}$$

$$A_3P(w') = \{abab\},$$

and hence $A_2P(w) \cap A_3P(w') \neq \emptyset$, but $c \in \Sigma_c$ satisfies $wc \in L \setminus K$ and $w'c \in K$. Thus K is robust with respect to symbol replacements but vulnerable to a combination of replacements and removals. □

B. Existence of supervisors

Our objective in this subsection is to provide a necessary and sufficient condition for the existence of a supervisor that achieves a specification language in the presence of output corruption. To this end, we first introduce supervisors for an attack set and define controlled languages under attacks.

A *P-supervisor for a language $L \subset \Sigma^*$ and an attack set $\mathcal{A}$* is a function $f : \bigcup_{A\in\mathcal{A}} AP(L) \to 2^{\Sigma}$, where $AP(L)$ is the set of all possible output strings under the attack A, that is,

$$AP(L) := \big\{y \in \Delta^* : \exists w \in L \text{ s.t. } y \in AP(w)\big\}.$$

We will say that a supervisor f is *valid* if $f(w) \supset \Sigma_u$ for all $w \in \bigcup_{A\in\mathcal{A}} AP(L)$.

Given a P-supervisor f for a language L and an attack set $\mathcal{A}$, the *maximal language $L^{\max}_{f,A}$ controlled by f under the attack $A \in \mathcal{A}$* is defined inductively by $\epsilon \in L^{\max}_{f,A}$ and

$$w\sigma \in L^{\max}_{f,A} \quad \Leftrightarrow \quad w \in L^{\max}_{f,A}, \quad w\sigma \in L, \quad \exists y \in AP(w) \text{ s.t. } \sigma \in f(y),$$

whereas the *minimal language $L^{\min}_{f,A}$ ($\subset L^{\max}_{f,A}$) controlled by f under the attack $A \in \mathcal{A}$* is defined inductively by $\epsilon \in L^{\min}_{f,A}$ and

$$w\sigma \in L^{\min}_{f,A} \quad \Leftrightarrow \quad w \in L^{\min}_{f,A}, \quad w\sigma \in L, \quad \forall y \in AP(w),\ \sigma \in f(y).$$

In general, $L^{\min}_{f,A} \subset L^{\max}_{f,A}$, but in the absence of an attack, i.e., $A = A_{\text{id}}$, both languages coincide. By construction, $L^{\max}_{f,A}$ and $L^{\min}_{f,A}$ are prefix closed.

By definition, $L^{\max}_{f,A}$ is the largest language that the attack A could enforce by exposing to the supervisor f an appropriate corrupted output string in $AP(w)$. On the other hand, $L^{\min}_{f,A}$ is the smallest language that the attack A could enforce. In other words, A could not reject any string in this set by the choice of corrupted output strings in $AP(w)$.

The following result provides a necessary and sufficient condition on a language $K \subset L$ for the existence of a P-supervisor f for an attack set $\mathcal{A}$ whose minimal and maximal controlled languages $L^{\min}_{f,A}$, $L^{\max}_{f,A}$ are both equal to K.

Theorem 4: For every nonempty prefix-closed set $K \subset L$ and every attack set $\mathcal{A}$:

1) There exists a valid P-supervisor f for $\mathcal{A}$ such that $L_{f,A}^{\min} = L_{f,A}^{\max} = K$ for all $A \in \mathcal{A}$ if and only if K is controllable and P-observable under $\mathcal{A}$;

2) If K is controllable and P-observable under $\mathcal{A}$, then the following map $f : \bigcup_{A\in\mathcal{A}} AP(L) \to 2^{\Sigma}$ defines a valid P-supervisor for which $L_{f,A}^{\min} = L_{f,A}^{\max} = K$ for all $A \in \mathcal{A}$:

$$f(y) := \Sigma_u \cup \left\{\sigma \in \Sigma_c : \exists w \in K, A \in \mathcal{A} \text{ s.t. } [y \in AP(w),\ w\sigma \in K]\right\},\ \forall y \in \bigcup_{A\in\mathcal{A}} AP(L). \tag{6}$$

When the set of attacks $\mathcal{A}$ contains only A_{id}, Theorem 4 specializes to the case without attacks.

Remark 5: As the proof below shows, in order to obtain controllability in item 1), it is enough that $L_{f,A}^{\max} = K$ or $L_{f,A}^{\min} = K$ for some $A \in \mathcal{A}$. □

Proof of Theorem 4: We first prove that $L_{f,A}^{\max} = K$ for some $A \in \mathcal{A}$ implies the controllability of K. Pick some word $\bar{w} \in K\Sigma_u \cap L$. Such a word must be of the form $\bar{w} = w\sigma \in L$ such that $w \in K$ and $\sigma \in \Sigma_u$. The supervisor f is defined for all the strings y that are produced by an attacker. Therefore, if $w \in K \subset L$, then $f(y)$ is defined for every $y \in AP(w)$. Since f is a valid supervisor, any uncontrollable event belongs to $f(y)$, in particular, $\sigma \in f(y)$. Since $w \in K = L_{f,A}^{\max}$, it now follows by the definition of $L_{f,A}^{\max}$ that $w\sigma \in L_{f,A}^{\max} = K$, which shows that $K\Sigma_u \cap L \subset K$, and therefore K is controllable. From $L_{f,A}^{\min} = K$ instead of $L_{f,A}^{\max} = K$, controllability can be obtained in the same way.

Next we prove that K is P-observable under an attack set $\mathcal{A}$ by using the fact that $K = L_{f,A}^{\min} = L_{f,A}^{\max}$ for all $A \in \mathcal{A}$. To do so, we employ the statement (4), which is equivalent to observability under attacks. Pick a pair of words $w, w' \in K$ such that

$$\exists A, A' \text{ s.t. } AP(w) \cap A'P(w') \neq \emptyset,$$

and an arbitrary symbol $\sigma \in \Sigma$ such that $w\sigma, w'\sigma \in L$. If such a symbol σ does not exist, then all σ satisfy $w\sigma \notin L$ or $w'\sigma \notin L$, and we immediately conclude that K is observable under $\mathcal{A}$ from (4). If such σ exists and $w\sigma \in K = L_{f,A}^{\min}$, then by the definition of $L_{f,A}^{\min}$, we must have

$$w \in L_{f,A}^{\min} = K, \quad w\sigma \in L, \quad \forall y \in AP(w),\ \sigma \in f(y).$$

Since $AP(w) \cap A'P(w') \neq \emptyset$, we must then have

$$w' \in L_{f,A'}^{\max} = K, \quad w'\sigma \in L, \quad \exists y \in A'P(w') \text{ s.t. } \sigma \in f(y).$$

Consequently, $w'\sigma \in L^{\max}_{f,A'} = K$. Alternatively, if $w\sigma \in L \setminus K$, then $w\sigma \notin K = L^{\max}_{f,A}$ and we must have

$$w \in L^{\max}_{f,A} = K, \quad w\sigma \in L, \quad \forall y \in AP(w),\ \sigma \notin f(y).$$

Since $AP(w) \cap A'P(w') \neq \emptyset$, we must then have

$$w' \in L^{\min}_{f,A'} = K, \quad w'\sigma \in L, \quad \exists y \in A'P(w') \text{ s.t. } \sigma \notin f(y),$$

and hence $w'\sigma \notin L^{\min}_{f,A'} = K$. This shows that (4) holds, and therefore K is P-observable under $\mathcal{A}$.

To prove the existence of the supervisor in item 1 (and also the statement in item 2), pick the supervisor f according to (6). We prove by induction on the word length that the supervisor f so defined satisfies $K = L^{\max}_{f,A} = L^{\min}_{f,A}$ for all $A \in \mathcal{A}$. The basis of induction is the empty string ϵ that belongs to $L^{\max}_{f,A}$ and $L^{\min}_{f,A}$ because of the definition of these sets and belongs to K because this set is prefix-closed.

Suppose now that K, $L^{\max}_{f,A}$, and $L^{\min}_{f,A}$ have exactly the same words of length $n \geq 0$, and pick a word $w'\sigma \in L^{\max}_{f,A}$ of length $n+1$ with $\sigma \neq \epsilon$. We show that $w'\sigma \in K$ as follows. Since $w' \in L^{\max}_{f,A}$ has length n, we know by the induction hypothesis that $w' \in K$. On the other hand, since $w'\sigma \in L^{\max}_{f,A}$, we must have

$$w' \in L^{\max}_{f,A}, \quad w'\sigma \in L, \quad \exists y \in AP(w') \text{ s.t. } \sigma \in f(y).$$

If $\sigma \in \Sigma_u$, then we see from controllability that $w'\sigma \in K$. Let us next consider the case $\sigma \in \Sigma_c$. By the definition (6) of f, $\sigma \in f(y)$ must mean that

$$\exists w \in K,\ \bar{A} \in \mathcal{A} \quad \text{s.t.} \quad [y \in \bar{A}P(w),\ w\sigma \in K]. \tag{7}$$

We therefore have

$$w, w' \in K, \quad AP(w') \cap \bar{A}P(w) \neq \emptyset, \quad w\sigma \in K, \quad w'\sigma \in L.$$

Since K is controllable and observable under $\mathcal{A}$, Proposition 3 shows that $w'\sigma \in K$. This shows that any word of length $n+1$ in $L^{\max}_{f,A}$ also belongs to K. Since $L^{\min}_{f,A} \subset L^{\max}_{f,A}$, it follows that any word of length $n+1$ in $L^{\min}_{f,A}$ also belongs to K.

Conversely, for a word $w'\sigma \in K \subset L$ of length $n+1$ with $\sigma \neq \epsilon$, we prove $w'\sigma \in L^{\min}_{f,A}$ as follows. Since K is prefix closed, we have $w' \in K$. The induction hypothesis shows that $w' \in L^{\min}_{f,A}$. To obtain $w'\sigma \in L^{\min}_{f,A}$, we need to show that

$$w' \in L^{\min}_{f,A}, \quad w'\sigma \in L, \quad \forall y \in AP(w'),\ \sigma \in f(y).$$

The first statement is a consequence of the induction hypothesis (as discussed above). The second statement is a consequence of the fact that $w'\sigma \in K \subset L$. As regards the third statement, if $\sigma \in \Sigma_u$, then $\sigma \in f(y)$ for all $y \in AP(w')$ by definition. It is therefore enough to show that $\sigma \in \Sigma_c$ leads to $\sigma \in f(y)$, that is, the statement (7), for every $y \in AP(w')$. We obtain (7) for the particular case $\bar{A} = A$, $w = w' \in K$. Thus any word of length $n+1$ in K also belongs to $L^{\min}_{f,A}$ and hence also to $L^{\max}_{f,A} \supset L^{\min}_{f,A}$, which completes the induction step. ■

C. *Observability under insertion-removal attacks*

In this subsection, we consider attacks that insert and remove certain symbols from output strings, and reduce observability under such attacks to the conventional observability in the non-attacked case.

Given a set of symbols $\alpha \subset \Delta$ in the observation alphabet, we define the *insertion-removal attack* $A_\alpha : \Delta^* \to 2^{\Delta^*}$ that maps each string $u \in \Delta^*$ to the set of all strings $v \in \Delta^*$ that can be obtained from u by an arbitrary number of insertions or removals of symbols in α. We say that A_α corresponds to an *attack on the output symbols in* α. In this context, it is convenient to also define the corresponding *α-removal observation map* $R_{\neg\alpha} : \Delta \to (\Delta \cup \{\epsilon\})$ by

$$R_{\neg\alpha}(t) = \begin{cases} \epsilon & t \in \alpha \\ t & t \notin \alpha. \end{cases} \tag{8}$$

The α-removal observation map can be extended to strings of events in the same way as the observation map P. This α-removal observation map allows us to define the attack $A_\alpha : \Delta^* \to 2^{\Delta^*}$ on the output symbol in α as follows:

$$A_\alpha(u) = \big\{v \in \Delta^* : R_{\neg\alpha}(u) = R_{\neg\alpha}(v)\big\}. \tag{9}$$

Example 6: Let $\alpha = \{t_1\} \subset \{t_1, t_2\}$. Then $R_{\neg\alpha}(t_1t_2) = t_2$, and

$$A_\alpha(t_1t_2) = \{t_1^n t_2 t_1^m : n, m \geq 0\} = \big\{v \in \Delta^* : R_{\neg\alpha}(t_1t_2) = R_{\neg\alpha}(v)\big\}. \qquad \square$$

The next result shows that the observability in Definition 2 under insertion-removal attacks is equivalent to the usual observability (without attacks) for an appropriate set of output maps. Note that the composition $R_{\neg\alpha} \circ P : \Delta \to (\Delta \cup \{\epsilon\})$ can be regarded as an observation map (in the usual sense, i.e., without attacks).

Theorem 7: For every nonempty prefix-closed language $K \subset L$ and insertion-removal attack set $\mathcal{A} = \{A_{\alpha_1}, A_{\alpha_2}, \ldots, A_{\alpha_M}\}$ consisting of $M \geq 1$ observation attacks, K is P-observable under the set of attacks $\mathcal{A}$ if and only if K is $(R_{\neg\alpha} \circ P)$-observable for every set $\alpha := \alpha_i \cup \alpha_j$, $\forall i, j \in \{1, 2, \ldots, M\}$.

Theorem 7 implies that one can use the standard test for DES observability (without attacks) in [32] to determine observability under insertion-removal attacks (Definition 2).

Remark 8 (Computational complexity of observability test for insertion-removal attacks): Consider a language $L = L(G)$ generated by a finite automaton $G = (X, \Sigma, \xi, x_0)$ and a specification language $K = L(G_K) \subset L$ generated by a finite automaton $G_K = (R, \Sigma, \eta, r_0)$. According to Thereom 7, in order to test observability under an insertion-removal attack set $\mathcal{A}$, it is enough to construct $|\mathcal{A}|^2$ test automatons, each of which verifies the usual observability. Since the computational complexity of verifying the usual observability with the test automaton is $O(|X| \cdot |R|^2 \cdot |\Sigma_c|)$ (see, e.g., [34, Sec. 3.7]), the total complexity for this observability test under attacks is $O(|X| \cdot |R|^2 \cdot |\Sigma_c| \cdot |\mathcal{A}|^2)$. □

The following result is the key step in proving Theorem 7.

Lemma 9: Given any two sets $\alpha_1, \alpha_2 \subset \Delta$ and $\alpha := \alpha_1 \cup \alpha_2$, we have that

$$(v, v') \in \ker R_{\neg\alpha} \quad \Rightarrow \quad A_{\alpha_1}(v) \cap A_{\alpha_2}(v') \neq \emptyset, \tag{10}$$

where $\ker R_{\neg\alpha}$ is defined by

$$\ker R_{\neg\alpha} := \big\{(v, v') \in \Delta^* \times \Delta^* : R_{\neg\alpha}(v) = R_{\neg\alpha}(v')\big\}. \tag{11}$$

Proof: To prove this result, we must show that given two strings $v, v' \in \Delta^*$ such that $R_{\neg\alpha}(v) = R_{\neg\alpha}(v')$, there exists a third string $y \in \Delta^*$ that belongs both to $A_{\alpha_1}(v)$ and $A_{\alpha_2}(v')$. In view of (9), this means that $y \in \Delta^*$ must satisfy

$$R_{\neg\alpha_1}(v) = R_{\neg\alpha_1}(y), \quad R_{\neg\alpha_2}(v') = R_{\neg\alpha_2}(y). \tag{12}$$

The desired string y can be constructed through the following steps:

1) Start with the string $y_1 := R_{\neg\alpha_1}(v)$, which is obtained by removing from v all symbols in α_1. Since $\alpha := \alpha_1 \cup \alpha_2$, we have $R_{\neg\alpha} = R_{\neg\alpha_2} \circ R_{\neg\alpha_1}$. Hence

$$R_{\neg\alpha_2}(y_1) = R_{\neg\alpha_2}(R_{\neg\alpha_1}(v)) = R_{\neg\alpha}(v) = R_{\neg\alpha}(v'),$$

and the strings y_1 and v' still only differ by symbols in α_1 and α_2.

2) Construct y_2 by adding to y_1 suitable symbols in α_1 so that $R_{\neg\alpha_2}(y_2) = R_{\neg\alpha_2}(v')$. This is possible because y_1 and v' only differ by symbols in α_1 and α_2 and y_1 has no symbols in α_1 by definition. To get $R_{\neg\alpha_2}(y_2) = R_{\neg\alpha_2}(v')$, we do not care about the symbols in α_2, so we just have to insert into y_1 the symbols in α_1 that appear in v' (at the appropriate locations).

3) By construction,

$$R_{\neg\alpha_1}(y_2) = y_1 = R_{\neg\alpha_1}(v),$$

and hence the original v and y_2 only differ by symbols in α_1. Moreover, $R_{\neg\alpha_2}(y_2) = R_{\neg\alpha_2}(v')$, that is, v' and y_2 only differ by symbols in α_2. We therefore conclude that $y := y_2$ satisfies (12) and hence belongs to both $A_{\alpha_1}(v)$ and $A_{\alpha_2}(v')$. ■

Proof of Theorem 7: By definition, K is P-observable under the set of attacks $\mathcal{A}$ if and only if

$$R_{A_{\alpha_i},A_{\alpha_j}} \subset \mathrm{act}_{K\subset L}, \qquad \forall i,j \in \{1,2,\ldots,M\}.$$

Also, by definition, K is $(R_{\neg\alpha} \circ P)$-observable for the set $\alpha := \alpha_i \cup \alpha_j$ if and only if

$$\ker(R_{\neg\alpha} \circ P) \subset \mathrm{act}_{K\subset L}\,.$$

To prove the result, it therefore suffices to show that, for all $i,j \in \{1,2,\ldots,M\}$, we have that

$$R_{A_{\alpha_i},A_{\alpha_j}} = \ker(R_{\neg\alpha} \circ P), \qquad \alpha := \alpha_i \cup \alpha_j.$$

To show that this equality holds, first pick a pair $(w,w') \in R_{A_{\alpha_i},A_{\alpha_j}}$, which means by the definition (2) of $R_{A_{\alpha_i},A_{\alpha_j}}$ that there exists a string $y \in \Delta^*$ that belongs both to $A_{\alpha_i}P(w)$ and $A_{\alpha_j}P(w')$, and therefore

$$\begin{aligned} y \in A_{\alpha_i}P(w) \quad &\Leftrightarrow \quad R_{\neg\alpha_i}\big(P(w)\big) = R_{\neg\alpha_i}(y) \\ y \in A_{\alpha_j}P(w') \quad &\Leftrightarrow \quad R_{\neg\alpha_j}\big(P(w')\big) = R_{\neg\alpha_j}(y) \end{aligned}$$

Since $\alpha := \alpha_i \cup \alpha_j$, we have that $R_{\neg\alpha} = R_{\neg\alpha_j} \circ R_{\neg\alpha_i} = R_{\neg\alpha_i} \circ R_{\neg\alpha_j}$, and consequently

$$\begin{aligned} R_{\neg\alpha}\big(P(w)\big) = R_{\neg\alpha_j}\big(R_{\neg\alpha_i}\big(P(w)\big)\big) = R_{\neg\alpha_j}\big(R_{\neg\alpha_i}(y)\big) \\ = R_{\neg\alpha_i}\big(R_{\neg\alpha_j}(y)\big) = R_{\neg\alpha_i}\big(R_{\neg\alpha_j}\big(P(w')\big)\big) = R_{\neg\alpha}\big(P(w')\big). \end{aligned}$$

Hence $(w,w') \in \ker(R_{\neg\alpha} \circ P)$. We have thus shown that $R_{A_{\alpha_i},A_{\alpha_j}} \subset \ker(R_{\neg\alpha} \circ P)$.

To prove the reverse inclusion, pick a pair $(w, w') \in \ker(R_{\neg\alpha} \circ P)$, which means that $R_{\neg\alpha}\big(P(w)\big) = R_{\neg\alpha}\big(P(w')\big)$, and therefore $\big(P(w), P(w')\big) \in \ker R_{\neg\alpha}$ by the definition (11) of $\ker R_{\neg\alpha}$. In conjunction with Lemma 9, this leads to

$$A_{\alpha_i}\big(P(w)\big) \cap A_{\alpha_j}\big(P(w')\big) \neq \emptyset.$$

Therefore we have $(w, w') \in R_{A_{\alpha_i}, A_{\alpha_j}}$. This shows that $\ker(R_{\neg\alpha} \circ P) \subset R_{A_{\alpha_i}, A_{\alpha_j}}$, which concludes the proof. ■

III. Realization of Supervisors under Attacks

The objective of this section is to describe the supervisor f in (6) through an automaton for replacement-removal attacks, which will be introduced in Section III A and are a special class of the attacks considered in Section II. In Section III B, we provide a construction for an automaton that describes a supervisor for replacement-removal attacks and then formally prove in Section III C that the proposed automaton represents the supervisor f in (6).

A. Replacement-removal attack

The results in this section are specific to attacks that consist of replacement and/or removal of specific symbols from the output strings. For a given replacement-removal map $\phi : \Delta \to 2^{\Delta \cup \{\epsilon\}}$ that maps each output symbol to a (possibly empty) set of symbols, the corresponding attack map $A : \Delta^* \to 2^{\Delta^*}$ is defined by $A(\epsilon) = \{\epsilon\}$ and

$$A(yt) := \big\{w_y w_t : w_y \in A(y), w_t \in \phi(t)\big\} \tag{13}$$

for all $y \in \Delta^*$ and $t \in \Delta$. Sets of attacks of this form are called *replacement-removal attack sets*. Recalling that $AP := A \circ P$, we conclude from (13) that

$$AP(s\sigma) = AP(s)AP(\sigma), \qquad \forall s \in \Sigma^*,\ \sigma \in \Sigma. \tag{14}$$

In what follows, it will be convenient to define the inverse map AP^{-1} by

$$AP^{-1}(y) := \big\{w \in \Sigma^* : y \in AP(w)\big\}, \qquad \forall y \in \Delta^*,$$

and extend it to languages $L_\Delta \subset \Delta^*$:

$$AP^{-1}(L_\Delta) := \big\{w \in \Sigma^* : \exists y \in L_\Delta \text{ s.t. } y \in AP(w)\big\}.$$

The following lemma provides basic properties of AP and AP^{-1} for a replacement-removal attack A:

Lemma 10: Consider a replacement-removal attack A.

1) *For all $w, s \in \Sigma^*$, we have*

$$AP(ws) = AP(w)AP(s). \tag{15}$$

2) *Let $y_1, \ldots, y_m \in \Delta$ and $w_1, \ldots, w_n \in \Sigma$. If $y_1 \cdots y_m \in AP(w_1 \cdots w_n)$, then there exist $i_1, \ldots, i_m$ such that $1 \le i_1 < i_2 < \cdots < i_m \le n$ and*

$$\begin{aligned} &\epsilon \in AP(w_1 \cdots w_{i_1-1}), \quad y_1 \in AP(w_{i_1}) \\ &\epsilon \in AP(w_{i_1+1} \cdots w_{i_2-1}), \quad y_2 \in AP(w_{i_2}) \\ &\quad \vdots \\ &\epsilon \in AP(w_{i_m+1} \cdots w_n), \end{aligned}$$

where, for example, if $i_2 - i_1 = 1$, then the condition $\epsilon \in AP(w_{i_1} \cdots w_{i_2-1})$ is omitted.

3) *For all $y, v \in \Delta^*$, we have*

$$AP^{-1}(yv) = AP^{-1}(y)AP^{-1}(v). \tag{16}$$

Proof: 1) If $s = \epsilon$, then (15) holds for all $w \in \Sigma^*$ because $AP(\epsilon) = \{\epsilon\}$. For $s = s_1 \cdots s_k$ with $s_i \in \Sigma$, we obtain (15), by using (14) iteratively. Thus we have (15) holds for all $w, s \in \Sigma^*$.

2) We prove the item 2 by induction on the word length of $y_1 \cdots y_m$. In the case $m = 1$, suppose that

$$y_1 \in AP(w_1 \cdots w_n) = AP(w_1) \cdots AP(w_n).$$

If $y_1 \notin AP(w_i)$ for every $i = 1, \ldots, n$, then $y_1 \notin AP(w_1 \cdots w_n)$. Hence $y_1 \in AP(w_i)$ for at least one i. Let $y_1 \in AP(w_i)$ hold for $i = j_1, \ldots, j_k$. It follows that there exists $l \in \{j_1, \ldots, j_k\}$ such that

$$\epsilon \in AP(w_1 \cdots w_{l-1}),\ y_1 \in AP(w_l),\ \epsilon \in AP(w_{l+1} \cdots w_n).$$

The desired i_1 is such an index l.

Suppose now that the item 2 holds with $y_1 \cdots y_k$ of length $k \ge 1$ and that

$$y_1 \cdots y_k y_{k+1} \in AP(w_1 \cdots w_n) = AP(w_1) \cdots AP(w_n).$$

Similarly to the case $m = 1$, we see that there exists i_{k+1} such that

$$y_1 \cdots y_k \in AP(w_1 \cdots w_{i_{k+1}-1}), \quad y_{k+1} \in AP(w_{i_{k+1}}), \quad \epsilon \in AP(w_{i_{k+1}+1} \cdots w_n).$$

Applying the induction hypothesis to $y_1 \cdots y_k \in AP(w_1 \cdots w_{i_{k+1}-1})$, we have the desired conclusion with $y_1 \cdots y_{k+1}$ of length $k+1$.

3) If (16) holds for all $y \in \Delta^*$ and $v \in \Delta \cup \{\epsilon\}$, then an iterative calculation shows that (16) holds for all $y \in \Delta^*$ and all $v \in \Delta^*$. Hence it is enough to prove that (16) holds for all $y \in \Delta^*$ and $t \in \Delta \cup \{\epsilon\}$.

Suppose that $w \in AP^{-1}(y)AP^{-1}(t)$ for $y \in \Delta^*$ and $t \in \Delta \cup \{\epsilon\}$. Then

$$\exists w_1 \in AP^{-1}(y),\ w_2 \in AP^{-1}(t)\ \text{ s.t. }\ w = w_1 w_2.$$

Since $y \in AP(w_1)$ and $t \in AP(w_2)$, it follows that

$$yt \in AP(w_1)AP(w_2) = AP(w),$$

which implies $w \in AP^{-1}(yt)$. We therefore have $AP^{-1}(y)AP^{-1}(t) \subset AP^{-1}(yt)$.

Next we prove the converse inclusion. If $t = \epsilon$, then we have from $AP(\epsilon) = \{\epsilon\}$ that $AP^{-1}(y)AP^{-1}(t) \supset AP^{-1}(yt)$.

Let us next consider the case $t \neq \epsilon$. Suppose that $w \in AP^{-1}(yt)$ for $y \in \Delta^*$ and $t \in \Delta$. Since $yt \in AP(w)$ and $yt \neq \epsilon$, it follows that $w \neq \epsilon$. Let $w = w_1 \cdots w_n$ with $w_i \in \Sigma$. From the item 2, there exists $i \in \{1, \ldots, n\}$ such that $y \in AP(w_1 \cdots w_i)$ and $t \in AP(w_{i+1} \cdots w_n)$. Hence if we define $w_1 := w_1 \cdots w_i$ and $w_2 := w_{i+1} \cdots w_n$, then

$$w = w_1 w_2,\ w_1 \in AP^{-1}(y),\ w_2 \in AP^{-1}(t).$$

Hence $w \in AP^{-1}(y)AP^{-1}(t)$. Thus $AP^{-1}(y)AP^{-1}(t) \supset AP^{-1}(yt)$, which completes the proof. ■

B. Supervisors described by observers

Inspired by the non-attacked results in [30], we first construct an observer that is resilient against replacement-removal attacks, which plays an important role in the representation of the supervisor f in (6).

Consider a specification automaton $G_K = (R, \Sigma, \eta, r_0)$ with a set Δ of output symbols and an output map $P : \Sigma \to (\Delta \cup \{\epsilon\})$. We define the *unobservable reach* $\mathrm{UR}_A(r)$ *of each state* $r \in R$ *under a replacement and removal attack* $A \in \mathcal{A}$ by

$$\mathrm{UR}_A(r) := \big\{ z \in R : \exists u \in AP^{-1}(\epsilon)\ \text{ s.t. }\ z = \eta(r, u) \big\},$$

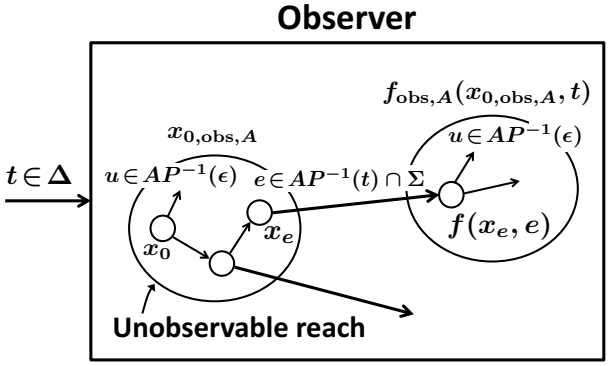

Fig. 3: Observer $\mathrm{Obs}_A(G_K)$ for an attack A.

which can be extended to a set of states $B \subset R$ by

$$\mathrm{UR}_A(B) := \bigcup_{r \in B} \mathrm{UR}_A(r).$$

To estimate the current state r of the specification automaton G_K in the presence of a replacement-removal attack $A \in \mathcal{A}$, the observer $\mathrm{Obs}_A(G_K) = (R_{\mathrm{obs},A}, \Delta, \eta_{\mathrm{obs},A}, r_{0,\mathrm{obs},A})$ is constructed in the following iterative procedure:

1) Define $r_{0,\mathrm{obs},A} := \mathrm{UR}_A(r_0) \subset R$ and set $R_{\mathrm{obs},A} = \{r_{0,\mathrm{obs},A}\}$.

2) For each set of states $B \in R_{\mathrm{obs},A}$ and $t \in \Delta$, if $\eta(r_e, e)$ is defined for some $r_e \in B$ and $e \in AP^{-1}(t) \cap \Sigma$, then define

$$\eta_{\mathrm{obs},A}(B, t) := \mathrm{UR}_A\big(\{r \in R : \exists r_e \in B,\ e \in AP(t)^{-1} \cap \Sigma \text{ s.t. } r = \eta(r_e, e)\}\big) \tag{17}$$

and add this set to $R_{\mathrm{obs},A}$; otherwise $\eta_{\mathrm{obs},A}(B, t)$ is not defined.

3) Go back to step 2) until no more sets can be added to $R_{\mathrm{obs},A}$.

Fig. 3 illustrates the observer $\mathrm{Obs}_A(G_K)$.

Example 1 (cont.): Consider the specification language $K = L(G_K)$ and the attack A_1 as in Example 1. Fig. 4 shows the observers obtained through the procedure described above for the attacks A_{id} (absence of attack) and A_1. The structure of the observer $\mathrm{Obs}_{A_1}(G_K)$ is similar to that of $\mathrm{Obs}_{A_{\mathrm{id}}}(G_K)$, but $\mathrm{Obs}_{A_1}(G_K)$ transitions by using only the length of the observed strings, which guarantees the robustness of K against replacement attacks. □

The next result provides the realization of the P-supervisor f in (6) for the language L and the specification K. This realization is base on a specification automaton $G_K := (R, \Sigma, \eta, r_0)$ with $L(G_K) = K$ and a family of observer automata $\mathrm{Obs}_A(G_K)$ designed using the procedure outlined above. Without loss of generality, we can restrict the domain of f to $\bigcup_{A \in \mathcal{A}} AP(K)$.

(a) Observer $\mathrm{Obs}_{A_{\mathrm{id}}}(G_K)$ for the attack A_{id} (non-attacked case).

(b) Observer $\mathrm{Obs}_{A_1}(G_K)$ for the attack A_1.

Fig. 4: Observers in Example 1.

Theorem 11: Consider a nonempty prefix-closed set $K \subset L$ and a replacement-removal attack set $\mathcal{A}$. Build a specification automaton $G_K := (R, \Sigma, \eta, r_0)$ with $L(G_K) = K$ and an observer $\mathrm{Obs}_A(G_K) := (R_{\mathrm{obs},A}, \Delta, \eta_{\mathrm{obs},A}, r_{0,\mathrm{obs},A})$ for every $A \in \mathcal{A}$. Define functions $\Psi : \bigcup_{A\in\mathcal{A}} R_{\mathrm{obs},A} \to 2^{\Sigma}$ and $\Phi_A : \bigcup_{A\in\mathcal{A}} AP(K) \to 2^{\Sigma}$ by

$$\Psi(r_{\mathrm{obs},A}) := \Sigma_u \cup \left\{\sigma \in \Sigma_c : \exists r \in r_{\mathrm{obs},A} \text{ s.t. } \eta(r,\sigma)!\right\} \tag{18}$$

$$\Phi_A(y) := \begin{cases} \Psi\big(\eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y)\big) & y \in AP(K) \\ \Sigma_u & y \notin AP(K). \end{cases} \tag{19}$$

Then the supervisor f in (6) *can be obtained using*

$$f(y) = \bigcup_{A\in\mathcal{A}} \Phi_A(y) \qquad \forall y \in \bigcup_{A\in\mathcal{A}} AP(K). \tag{20}$$

We can precompute the function $\Psi(r_{\mathrm{obs},A})$ for each observer state $r_{\mathrm{obs},A}$ and then can obtain the desired control action by looking at the current state $r_{\mathrm{obs},A}$ and the corresponding event set $\Psi(r_{\mathrm{obs},A})$ for all observers $\mathrm{Obs}_A(G_K)$ ($A \in \mathcal{A}$).

Remark 12: The branch $\Phi_A(y) = \Sigma_u$ for $y \notin AP(K)$ in (19) implies that once an attack A and a corrupted output string y satisfy $y \notin AP(K)$, the corresponding observer $\mathrm{Obs}_A(G_K)$ is not under operation because the actual attack must be different from the attack A that the observer assumes. Analogous to the case for linear time-invariant systems in [9]–[11], once $y \notin AP(K)$ is detected, the supervisor can exclude that attack from further consideration and stop updating the corresponding observer. □

C. Proof of Theorem 11

The following lemma shows that the state of $\mathrm{Obs}_A(G_K)$ is the set containing all the the states of G_K that could be reached from the initial state under the attack A.

Lemma 13: Consider a replacement-removal attack A. For all $y \in L(\mathrm{Obs}_A(G_K))$, we have

$$r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y) \quad \Leftrightarrow \quad \exists w \in AP^{-1}(y) \text{ s.t. } r = \eta(r_0, w). \tag{21}$$

Moreover, $L(\mathrm{Obs}_A(G_K)) = AP(L(G_K))$.

The proof of Lemma 13 relies on a key technical result that provides a more direct representation of $\eta_{\mathrm{obs},A}$ in (17) without unobservable reaches UR_A.

Lemma 14: Consider a replacement-removal attack A. Let $B \in R_{\mathrm{obs},A}$ and $t \in \Delta$. If $\eta(r_e, e)$ is defined for some $r_e \in B$ and $e \in AP^{-1}(t)$, then we define

$$\bar{\eta}_{\mathrm{obs},A}(B,t) := \big\{ r \in R : \exists r_e \in B,\ e \in AP^{-1}(t) \text{ s.t. } r = \eta(r_e, e) \big\}.$$

Then

$$\eta_{\mathrm{obs},A}(B,t)! \quad \Leftrightarrow \quad \bar{\eta}_{\mathrm{obs},A}(B,t)! \tag{22}$$

and

$$\eta_{\mathrm{obs},A}(B,t) = \bar{\eta}_{\mathrm{obs},A}(B,t). \tag{23}$$

Proof: First we prove (22). To prove this, it suffices to show that for all $B \in R_{\mathrm{obs},A}$ and $t \in \Delta$,

$$\begin{aligned} &\eta(r_e, e)! \text{ for some } r_e \in B \text{ and } e \in AP^{-1}(t) \cap \Sigma \\ \Leftrightarrow \quad &\eta(r_e, e)! \text{ for some } r_e \in B \text{ and } e \in AP^{-1}(t) \end{aligned} \tag{24}$$

By construction, ($\Rightarrow$) holds. We prove ($\Leftarrow$) as follows. Assume that there exist $\bar{r}_e \in B$ and $\bar{e} \in AP^{-1}(t)$ such that $\eta(\bar{r}_e, \bar{e})$ is defined. Since $t \in AP(\bar{e})$, Lemma 10 shows that there exist $e_1 \in \Sigma^*$, $e_2 \in \Sigma$, and $e_3 \in \Sigma^*$ such that

$$\bar{e} = e_1 e_2 e_3, \quad \epsilon \in AP(e_1), \quad t \in AP(e_2), \quad \epsilon \in AP(e_3). \tag{25}$$

Since $\eta(\bar{r}_e, \bar{e})!$, we also obtain $\eta(\eta(\bar{r}_e, e_1), e_2)!$. Therefore if we prove that

$$\bar{r}_e \in B, \quad \epsilon \in AP(e_1) \quad \Rightarrow \quad \eta(\bar{r}_e, e_1) \in B, \tag{26}$$

then ($\Leftarrow$) of (24) holds with $r_e = \eta(\bar{r}_e, e_1)$ and $e = e_2$.

Let us show (26). Since $B \in R_{\mathrm{obs},A}$, it follows that $B = \mathrm{UR}_A(\bar{B})$ for some $\bar{B} \subset R$. From $\bar{r}_e \in B$, we have $\bar{r}_e = \eta(r,u)$ for some $r \in \bar{B}$ and $u \in AP^{-1}(\epsilon)$. Since $\epsilon \in AP(e_1)$, it follows that $ue_1 \in AP^{-1}(\epsilon)$. Thus

$$\eta(\bar{r}_e, e_1) = \eta(r, ue_1) \in \mathrm{UR}_A(\bar{B}) = B,$$

which completes the proof of (22).

Next we prove (23). To show $\eta_{\mathrm{obs},A}(B,t) \subset \bar{\eta}_{\mathrm{obs},A}(B,t)$, suppose that $z \in \eta_{\mathrm{obs},A}(B,t)$. Then there exist

$$r \in \big\{r' \in R : \exists r_e \in B,\ e \in AP(t)^{-1} \cap \Sigma \text{ s.t. } r' = \eta(r_e, e)\big\} =: M$$

and $u \in AP^{-1}(\epsilon)$ such that $z = \eta(r,u)$. This implies that

$$\exists r_e \in B,\ e \in AP^{-1}(t) \cap \Sigma,\ u \in AP^{-1}(\epsilon) \text{ s.t. } z = \eta(r_e, eu).$$

Since $eu \in AP^{-1}(t)$, it follows that $z \in \bar{\eta}_{\mathrm{obs},A}(B,t)$. Thus $\eta_{\mathrm{obs},A}(B,t) \subset \bar{\eta}_{\mathrm{obs},A}(B,t)$

To prove the reverse inclusion, pick $z \in \bar{\eta}_{\mathrm{obs},A}(B,t)$. Then

$$\exists \bar{r}_e \in B,\ \bar{e} \in AP^{-1}(t) \text{ s.t. } z = \eta(\bar{r}_e, \bar{e}).$$

Since $t \in AP(\bar{e})$, it follows from Lemma 10 that there exist $e_1 \in \Sigma^*$, $e_2 \in \Sigma$, and $e_3 \in \Sigma^*$ such that (25) holds. Furthermore, as discussed above, $\bar{r}_e \in B$ and $\epsilon \in AP(e_1)$ lead to $\eta(\bar{r}_e, e_1) \in B$. Since $e_2 \in AP^{-1}(t) \cap \Sigma$, we have

$$\eta(\bar{r}_e, e_1 e_2) = \eta\big(\eta(\bar{r}_e, e_1), e_2\big) \in M.$$

Finally, since $\epsilon \in AP(e_3)$, it follows that

$$z = \eta(r_e, e_1 e_2 e_3) \in \mathrm{UR}_A(M) = \eta_{\mathrm{obs},A}(B,t).$$

Thus $\eta_{\mathrm{obs},A}(B,t) \supset \bar{\eta}_{\mathrm{obs},A}(B,t)$. This completes the proof. ■

We are now ready to prove Lemma 13.

Proof of Lemma 13: We prove (21) by induction on the word length. Since

$$r_{0,\mathrm{obs},A} = \big\{z \in R : \exists u \in AP^{-1}(\epsilon) \text{ s.t. } z = \eta(r_0, u)\big\},$$

we have that if $y = \epsilon$, then

$$r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y) \ \Leftrightarrow\ \exists u \in AP^{-1}(\epsilon) \text{ s.t. } r = \eta(r_0, u),$$

which means that (21) holds for $y = \epsilon$.

Suppose now that (21) holds for all words $y \in L\big(\mathrm{Obs}_A(G_K)\big)$ of length $n \geq 0$. Pick a word $y't \in L\big(\mathrm{Obs}_A(G_K)\big)$ of length $n+1$, where $y' \in L\big(\mathrm{Obs}_A(G_K)\big)$ has length n. By Lemma 14,

$$\begin{aligned}\eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y't) &= \eta_{\mathrm{obs},A}\big(\eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y'), t\big) \\ &= \bar{\eta}_{\mathrm{obs},A}\big(\eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y'), t\big)\end{aligned}$$

We therefore have

$$r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y't) \quad \Leftrightarrow \quad \exists r_e \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y'),\ e \in AP^{-1}(t) \ \text{ s.t. } \ r = \eta(r_e, e). \quad (27)$$

Since y' has length n, we know by the induction hypothesis that $r_e \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y')$ if and only if

$$\exists w' \in AP^{-1}(y') \ \text{ s.t. } \ r_e = \eta(r_0, w').$$

Combining this with (27), we obtain

$$r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y't) \quad \Leftrightarrow \quad \exists w' \in AP^{-1}(y'),\ e \in AP^{-1}(t) \ \text{ s.t. } \ r = \eta(r_0, w'e). \quad (28)$$

To prove $(\Rightarrow)$ in (21), suppose that $r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y't)$. Then from (28) and Lemma 10, $w := w'e$ satisfies

$$w \in AP^{-1}(y')AP^{-1}(t) = AP^{-1}(y't)$$

and $r = \eta(r_0, w)$, which implies that $(\Rightarrow)$ in (21) holds.

Let us next show that $(\Leftarrow)$ in (21) holds. Conversely, suppose that

$$\exists w \in AP^{-1}(y't) \ \text{ s.t. } \ r = \eta(r_0, w).$$

Since $AP^{-1}(y't) = AP^{-1}(y')AP^{-1}(t)$ from Lemma 10, it follows that

$$\exists w' \in AP^{-1}(y'),\ e \in AP^{-1}(t) \ \text{ s.t. } \ w = w'e.$$

Since $r = \eta(r_0, w'e)$, (28) shows that $r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs}}, y't)$, and hence we have $(\Leftarrow)$ in (21).

Next we prove $L\big(\mathrm{Obs}_A(G_K)\big) = AP\big(L(G_K)\big)$ by induction on the word length. The basis of induction is the empty string ϵ that belongs to $L\big(\mathrm{Obs}_A(G_K)\big)$ by definition, and belongs to $AP\big(L(G_K)\big)$ because $AP(\epsilon) = \{\epsilon\}$.

Suppose now that $L\big(\mathrm{Obs}_A(G_K)\big)$ and $AP\big(L(G_K)\big)$ have exactly the same words of length $n \geq 0$, and pick a word $y't \in \Delta^*$ of length $n+1$, where y' and t have length n and 1, respectively. Lemma 14 shows

$$\begin{aligned} & y't \in L\big(\mathrm{Obs}_A(G_K)\big) \\ & \quad \Leftrightarrow \quad \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y't)! \\ & \quad \Leftrightarrow \quad \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y')! \ \wedge \ \big[\exists r_e \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y'), e \in AP^{-1}(t) \ \text{ s.t. } \ \eta(r_e, e)!\big]. \end{aligned} \tag{29}$$

Suppose that $y't \in L(\mathrm{Obs}_A(G_K))$. We do not use the induction hypothesis to prove that $y't \in AP(L(G_K))$ holds. Since $y' \in L(\mathrm{Obs}_A(G_K))$ as well, it follows from (21) that

$$r_e \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y') \quad \Leftrightarrow \quad \exists w' \in AP^{-1}(y') \ \text{ s.t. } \ r_e = \eta(r_0, w'). \tag{30}$$

Combining this with (29), we have

$$\exists w' \in AP^{-1}(y'), \ e \in AP^{-1}(t) \ \text{ s.t. } \ \eta(r_0, w'e)!. \tag{31}$$

Hence $w := w'e$ satisfies $AP(w) = AP(w')AP(e) \ni y't$ and $w \in L(G_K)$, which implies $y't \in AP\big(L(G_K)\big)$. Thus we have $L\big(\mathrm{Obs}_A(G_K)\big) \subset AP\big(L(G_K)\big)$.

Conversely, suppose that $y't \in AP\big(L(G_K)\big)$. Then there exists $w \in L(G_K)$ such that $y't \in AP(w)$. Lemma 10 shows that

$$w = w_1 w_2, \ y' \in AP(w_1), \ t \in AP(w_2). \tag{32}$$

for some $w_1, w_2 \in \Sigma^*$. Since $w \in L(G_K)$ leads to $w_1 \in L(G_K)$, it follows that $y' \in AP\big(L(G_K)\big)$. Since y' has length n, the induction hypothesis gives $y' \in L\big(\mathrm{Obs}_A(G_K)\big)$. Hence we have $\eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y')!$ and (30). To prove $y't \in L\big(\mathrm{Obs}_A(G_K)\big)$, it suffices from (29) and (30) to show that (31) holds. Define $w' := w_1$ and $e := w_2$. Then we have $w' \in AP^{-1}(y')$ and $e \in AP^{-1}(t)$ from (32). Moreover,

$$w'e = w_1 w_2 = w \in L(G_K)$$

leads to $\eta(r_0, w'e)!$. We therefore obtain (31). Thus $L\big(\mathrm{Obs}_A(G_K)\big) \supset AP\big(L(G_K)\big)$, which completes the proof. ■

Using Lemma 13, we finally provide the proof of Theorem 11.

Proof of Theorem 11: Define $f_A : \bigcup_{A \in \mathcal{A}} AP(K) \to 2^{\Sigma}$ by

$$f_A(y) := \Sigma_u \cup \{\sigma \in \Sigma_c : \exists w \in K \ \text{ s.t. } \ y \in AP(w), \ w\sigma \in K\}$$

for all $y \in \bigcup_{A\in\mathcal{A}} AP(K)$. Then the supervisor f in (6) satisfies

$$f(y) = \bigcup_{A\in\mathcal{A}} f_A(y) \qquad \forall y \in \bigcup_{A\in\mathcal{A}} AP(K),$$

Hence, in order to obtain (20), we need to prove that $f_A = \Phi_A$ for each $A \in \mathcal{A}$.

By definition, $f_A(y) = \Sigma_u = \Phi_A(y)$ for all $y \notin AP(K)$. Suppose that $y \in AP(K)$. Since $L(G_K) = K$, it follows from Lemma 13 that

$$AP(K) = AP\big(L(G_K)\big) = L\big(\mathrm{Obs}_A(G_K)\big).$$

Since $y \in AP(K) = L\big(\mathrm{Obs}_A(G_K)\big)$, we conclude from (21) that

$$\exists r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y) \ \text{ s.t. } \ \eta(r,\sigma)! \quad \Leftrightarrow \quad \exists w \in AP^{-1}(y) \ \text{ s.t. } \ \eta(r_0, w\sigma)!.$$

Since $\eta(r_0, w\sigma)!$ means that $w, w\sigma \in K$, it follows that

$$\begin{aligned}
&\big\{\sigma \in \Sigma_c : \exists r \in \eta_{\mathrm{obs},A}(r_{0,\mathrm{obs},A}, y) \ \text{ s.t. } \ \eta(r,\sigma)!\big\} \\
&\quad = \{\sigma \in \Sigma_c : \exists w \in AP^{-1}(y) \ \text{ s.t. } \ w\sigma \in K\} \\
&\quad = \{\sigma \in \Sigma_c : \exists w \in K \ \text{ s.t. } \ y \in AP(w), \ w\sigma \in K\}
\end{aligned}$$

which implies that $f_A(y) = \Phi_A(y)$ also for every $y \in AP(K)$. This completes the proof. ■

IV. Test for Observability under Replacement-Removal Attacks

In this section, we propose an observability test under replacement-removal attacks in Section III A, inspired by product automata [32].

A. Product automata

Consider a language $L = L(G)$ generated by a finite automaton $G = (X, \Sigma, \xi, x_0)$. We want to test the observability under an attack set $\mathcal{A}$ of a specification language $K = L(G_K) \subset L$ generated by a finite automaton $G_K = (R, \Sigma, \eta, r_0)$.

For each $A, A' \in \mathcal{A}$, we construct a product automaton $T_{A,A'} = (Q, \Sigma_{T,A,A'}, \delta_{A,A'}, q_0)$ for the observability test, where $Q := R \times X \times R$, $q_0 := (r_0, x_0, r_0)$, and

$$\Sigma_{T,A,A'} := \big\{(\sigma,\sigma') \in (\Sigma \cup \{\epsilon\}) \times (\Sigma \cup \{\epsilon\}) \setminus \{(\epsilon,\epsilon)\} : AP(\sigma) \cap A'P(\sigma') \neq \emptyset\big\}.$$

The events in $\Sigma_{T_{A,A'}}$ are the pairs of events of the original automaton G that may not be distinguished under attacks A and A' by a supervisor.

The transition function $\delta_{A,A'} : Q \times \Sigma_{T,A,A'} \to Q$ is defined in the following way: For every $q = (r, x', r') \in Q$ and every $\sigma_T = (\sigma, \sigma') \in \Sigma_{T,A,A'}$, $\delta_{A,A'}(q, \sigma_T)$ is defined if (and only if) $\eta(r, \sigma)$, $\xi(x', \sigma')$, and $\eta(r', \sigma')$ are all defined. If $\delta_{A,A'}(q, \sigma_T)$ is defined, then

$$\delta_{A,A'}(q, \sigma_T) = \big(\eta(r, \sigma), \xi(x', \sigma'), \eta(r', \sigma')\big).$$

B. Observability test

For each $A, A' \in \mathcal{A}$, let us denote by $\mathrm{Ac}_{A,A'}(Q)$ the set of accessible states from the initial state q_0 by some string in $L(T_{A,A'})$, that is,

$$\mathrm{Ac}_{A,A'}(Q) := \big\{q \in Q : \exists (w, w') \in L(T_{A,A'}) \text{ s.t. } q = \delta_{A,A'}(q_0, (w, w'))\big\}.$$

The following theorem provides a necessary and sufficient condition for observability under replacement-removal attacks, which can be checked using a set of test automatons $\{T_{A,A'}\}_{A,A' \in \mathcal{A}}$.

Theorem 15: Consider a replacement-removal attack set $\mathcal{A}$ and the test automaton $T_{A,A'}$ defined above. Assume that a prefix-closed language $K \subset L(G)$ is controllable. Then K is not P-observable under $\mathcal{A}$ if and only if there exist $A, A' \in \mathcal{A}$, $(r, x', r') \in \mathrm{Ac}_{A,A'}(Q)$, and $\sigma \in \Sigma_c$ such that

$$\eta(r, \sigma)! \wedge \xi(x', \sigma)! \wedge \neg\eta(r', \sigma)!. \tag{33}$$

The proof of Theorem 15 is provided in the next subsection.

Remark 16 (Computational complexity of observability test for replacement-removal attacks): The number of the elements in $\mathrm{Ac}_{A,A'}(Q)$ is at most $|X| \cdot |R|^2$. Since we only have to check the state transition of the elements in $\mathrm{Ac}_{A,A'}(Q)$ driven by a controllable event, the total computational complexity to test observability is $O(|X| \cdot |R|^2 \cdot |\Sigma_c| \cdot |\mathcal{A}|^2)$ for the replacement-removal attack case. This complexity is the same as the one we derived in Remark 8 for insertion-removal attacks. □

Example 1 (cont.): Consider again the language $L(G)$, the specification language $K = L(G_K)$, and the attack set $\mathcal{A} = \{A_2, A_3\}$ in Example 1. To check observability under $\mathcal{A}$ by Theorem 15, we need to construct four test automatons T_{A_2,A_2}, T_{A_3,A_3}, T_{A_3,A_2}, and T_{A_2,A_3}. We see from T_{A_3,A_2} in Fig. 5 that K is not observable under $\mathcal{A}$. In fact, for the state (r_2, x_1', r_1') and the controllable event $c \in \Sigma_c$, we have

$$\eta(r_2, c)!, \ \xi(x_1', c)!, \text{ and } \neg\eta(r_1', c)!.$$

Thus K is not observable under $\mathcal{A}$, which is consistent with the discussion in Example 1. □

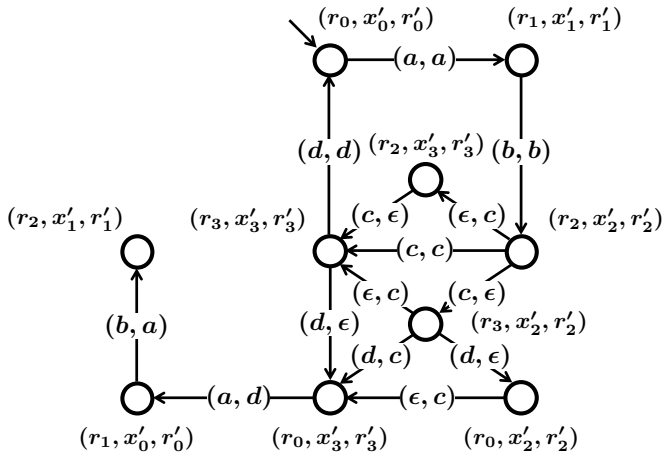

Fig. 5: Test automaton T_{A_3,A_2}.

C. The proof of Theorem 15

For all $(w, w') \in \Sigma^* \times \Sigma^*$ and all $(\sigma, \sigma') \in (\Sigma \cup \{\epsilon\}) \times (\Sigma \cup \{\epsilon\})$, we define

$$(w, w')(\sigma, \sigma') := (w\sigma, w'\sigma').$$

Using this notation, we define

$$\Sigma^*_{T,A,A'} := \big\{(\epsilon, \epsilon)\big\} \cap \big\{(\sigma_1, \sigma'_1) \cdots (\sigma_n, \sigma'_n) : (\sigma_i, \sigma'_i) \in \Sigma_{T,A,A'},\ \forall i = 1, \ldots, n,\ n \in \mathbb{N}\big\}.$$

Also, we define $\Sigma^0_{T,A,A'} := \{(\epsilon, \epsilon)\}$ and for each $n \geq 1$,

$$\Sigma^n_{T,A,A'} := \big\{(\sigma_1, \sigma'_1) \cdots (\sigma_n, \sigma'_n) : (\sigma_i, \sigma'_i) \in \Sigma_{T,A,A'},\ \forall i = 1, \ldots, n\big\}.$$

$\Sigma^n_{T,A,A'}$ is a subset of $\Sigma^*_{T,A,A'}$ whose elements have length $n \geq 0$.

The following proposition provides another representation of $\delta_{A,A'}(q_0, (w, w'))$ with initial states r_0 and x_0.

Lemma 17: Consider replacement-removal attacks A, A' and the test automaton $T_{A,A'}$ defined above. For every $(w, w') \in \Sigma^*_{T,A,A'}$, if $\delta_{A,A'}(q_0, (w, w'))$ is defined, then

$$\eta(r_0, w)!, \quad \xi(x_0, w')!, \quad \eta(r_0, w')! \tag{34}$$

$$\delta_{A,A'}\big(q_0, (w, w')\big) = \big(\eta(r_0, w), \xi(x_0, w'), \eta(r_0, w')\big). \tag{35}$$

Proof: We prove (34) and (35) by induction on the word length of (w, w'). The basis of induction is the empty string (ϵ, ϵ). Since $\delta_{A,A'}(q_0, (\epsilon, \epsilon)) = q_0 = (r_0, x_0, r_0)$ and $\eta(r_0, \epsilon) = r_0$, $\xi(x_0, \epsilon) = x_0$, $\eta(r_0, \epsilon) = r_0$, it follows that (34) and (35) hold.

Suppose now that for all $(w, w') \in \Sigma^n_{T,A,A'}$ with $n \geq 0$, if $\delta_{A,A'}\big(q_0, (w, w')\big)$ is defined, then (34) and (35) hold. Let $(\bar{w}, \bar{w}') \in \Sigma^{n+1}_{T,A,A'}$ satisfies $\delta_{A,A'}\big(q_0, (\bar{w}, \bar{w}')\big)$ is defined. Then there exist $(w, w') \in \Sigma^n_{T,A,A'}$ and $(\sigma, \sigma') \in \Sigma_{T,A,A'}$ such that $(\bar{w}, \bar{w}') = (w, w')(\sigma, \sigma')$ and

$$\delta_{A,A'}\big(q_0, (w, w')\big)!, \;\; \delta_{A,A'}\big(\delta_{A,A'}(q_0, (w, w')), (\sigma, \sigma')\big)!. \tag{36}$$

We know by the induction hypothesis that $\eta(r_0, w)$, $\xi(x_0, w')$, and $\eta(r_0, w')$ are defined and that $\delta_{A,A'}\big(q_0, (w, w')\big) = \big(\eta(r_0, w), \xi(x_0, w'), \eta(r_0, w')\big)$. Hence we see from the second statement of (36) that $\eta(r_0, \bar{w}) = \eta\big(\eta(r_0, w), \sigma\big)$ is defined. Similarly, $\xi(x_0, \bar{w}')$ and $\eta(r_0, \bar{w}')$ are defined. Furthermore, we have

$$\begin{aligned}\delta_{A,A'}\big(q_0, (\bar{w}, \bar{w}')\big) &= \delta_{A,A'}\big(\delta_{A,A'}(q_0, (w, w')), (\sigma, \sigma')\big) \\ &= \big(\eta(r_0, \bar{w}), \xi(x_0, \bar{w}'), \eta(r_0, \bar{w}')\big).\end{aligned}$$

Thus the desired statements (34) and (35) hold for all $(\bar{w}, \bar{w}') \in \Sigma^{n+1}_{T,A,A'}$, which completes the induction step. ■

The next lemma shows that the product automaton $T_{A,A'}$ tests state transitions by two events in K whose observation alphabets may not be distinguished under attacks A and A' by the supervisor:

Lemma 18: Consider replacement-removal attacks A, A'. For the test automaton $T_{A,A'}$ defined as in Section IV A, we have

$$\begin{aligned}L(T_{A,A'}) &:= \big\{(w, w') \in \Sigma^*_{T,A,A'} : \delta_{A,A'}(q_0, (w, w'))!\big\} \\ &= \big\{(w, w') \in \Sigma^* \times \Sigma^* : w, w' \in K, \;\; AP(w) \cap A'P(w') \neq \emptyset\big\} =: \hat{L}_{A,A'}.\end{aligned} \tag{37}$$

In the proof of Lemma 18, we use the lemma below that provides the property of $\Sigma^*_{T,A,A'}$ and a set equivalent to $\hat{L}_{A,A'}$.

Lemma 19: Consider two replacement-removal attacks A, A' and the test automaton $T_{A,A'}$ defined as in Section IV A. For all $(w, w') \in \Sigma^* \times \Sigma^*$, we have

$$(w, w') \in \Sigma^*_{T,A,A'} \quad \Leftrightarrow \quad AP(w) \cap A'P(w') \neq \emptyset. \tag{38}$$

Moreover, the language $\hat{L}_{A,A'}$ in (37) *satisfies*

$$\hat{L}_{A,A'} = \big\{(w, w') \in \Sigma^*_{T,A,A'} : w, w' \in K\big\}. \tag{39}$$

Proof: ($\Rightarrow$ of (38)) We show the proof by induction on the word length $(w, w') \in \Sigma^*_{T,A,A'}$. The basis of the induction, (ϵ, ϵ), satisfies $\epsilon \in AP(\epsilon) \cap A'P(\epsilon)$. We therefore have $AP(\epsilon) \cap A'P(\epsilon) \neq \emptyset$.

Suppose that all words $(w, w') \in \Sigma^n_{T,A,A'}$ satisfies $AP(w) \cap A'P(w) \neq \emptyset$. Let $(\bar{w}, \bar{w}') \in \Sigma^{n+1}_{T,A,A'}$. Then there exists $(w, w') \in \Sigma^n_{T,A,A'}$ and $(\sigma, \sigma') \in \Sigma_{T,A,A'}$ such that $\bar{w} = w\sigma$ and $\bar{w}' = w'\sigma'$. The induction hypothesis shows that $AP(w) \cap A'P(w') \neq \emptyset$, and by construction, we have $AP(\sigma) \cap A'P(\sigma') \neq \emptyset$. Thus $AP(\bar{w}) \cap A'P(\bar{w}') \neq \emptyset$.

($\Leftarrow$ of (38)) We split the proof into two cases:

$$[w = \epsilon \text{ or } w' = \epsilon] \quad \text{and} \quad [w \neq \epsilon \text{ and } w' \neq \epsilon].$$

First we consider the case [$w = \epsilon$ or $w' = \epsilon$]. Let $w = \epsilon$. Since $AP(\epsilon) \cap A'P(w') \neq \emptyset$ and since $AP(\epsilon) = \{\epsilon\}$, it follows that $\epsilon \in A'P(w')$. If $w' = \epsilon$, then $(w, w') = (\epsilon, \epsilon) \in \Sigma^*_{T,A,A'}$. Suppose that $w' \neq \epsilon$. Let $w' = w'_1 \cdots w'_k$ with $w'_1, \ldots, w'_k \in \Sigma$. Then

$$\epsilon \in AP(w') = AP(w'_1) \cdots AP(w'_k),$$

and hence $\epsilon \in AP(\epsilon) \cap AP(w'_i)$ for all $i = 1, \ldots, k$. Thus $(w, w') \in \Sigma^*_{T,A,A'}$.

Next we study the case [$w \neq \epsilon$ and $w' \neq \epsilon$]. Let $w = w_1 \cdots w_l$ with $w_1, \ldots, w_l \in \Sigma$ and $w' = w'_1 \cdots w'_k$ with $w'_1, \ldots, w'_k \in \Sigma$. Suppose that $AP(w) \cap A'P(w') = \{\epsilon\}$. Then

$$\epsilon \in AP(w) = AP(w_1) \cdots AP(w_l)$$
$$\epsilon \in AP(w') = AP(w'_1) \cdots AP(w'_k).$$

Hence $\epsilon \in AP(w_j)$ for all $j = 1, \ldots, l$ and $\epsilon \in AP(w'_i)$ for all $i = 1, \ldots, k$. Thus $(w, w') \in \Sigma^*_{T,A,A'}$.

Suppose that $AP(w) \cap A'P(w') \neq \{\epsilon\}$, then there exists $y \in \Delta^* \setminus \{\epsilon\}$ such that $y \in AP(w) \cap A'P(w')$. Let $y = y_1 \cdots y_m$ with $y_1, \ldots, y_m \in \Delta$. From Lemma 10, there exist $i_1, \ldots, i_m$ such that $y_1 \in AP(w_{i_1}), \ldots, y_m \in AP(w_{i_m})$ and

$$\epsilon \in AP(w_1 \cdots w_{i_1 - 1}) = AP(w_1) \cdots AP(w_{i_1 - 1})$$
$$\epsilon \in AP(w_{i_1 + 1} \cdots w_{i_2 - 1}) = AP(w_{i_1 + 1}) \cdots AP(w_{i_2 - 1})$$
$$\vdots$$
$$\epsilon \in AP(w_{i_m + 1} \cdots w_l) = AP(w_{i_m + 1}) \cdots AP(w_l),$$

which implies that $\epsilon \in AP(w_i)$ for all $i \neq i_1, \ldots, i_m$. We also have similar indices for w'. Thus $(w, w') \in \Sigma^*_{T,A,A'}$.

The second statement (39) directly follows from (38). ■

Using Lemmas 17 and 19, we prove Lemma 18.

Proof of Lemma 18: From Lemmas 17 and 19, we obtain $L(T_{A,A'}) \subset \hat{L}_{A,A'}$.

Let us prove $L(T_{A,A'}) \supset \hat{L}_{A,A'}$. From Lemma 19, it is enough to show this inclusion by induction on the word length $(w, w') \in \Sigma^*_{T,A,A'}$. Since $AP(\epsilon) \cap A'P(\epsilon) = \{\epsilon\}$, the basis of the induction, (ϵ, ϵ), belongs to both sets $L(T_{A,A'})$ and $\hat{L}_{A,A'}$.

Suppose that all words in $\hat{L}_{A,A'}$ whose length is $n \geq 0$ belong to $L(T_{A,A'})$. Let $(\bar{w}, \bar{w}') \in \hat{L}_{A,A'}$ have length $n+1$, that is, $(\bar{w}, \bar{w}')$ satisfy $\Sigma^{n+1}_{T,A,A'}$ and $\bar{w}, \bar{w}' \in K$. Then

$$(\bar{w}, \bar{w}') = (w, w')(\sigma, \sigma')$$

for some $(w, w') \in \Sigma^n_{T,A,A'}$ and some $(\sigma, \sigma') \in \Sigma_{T,A,A'}$. Since $w, w' \in K$, the induction hypothesis shows that $\eta(r_0, w)$, $\xi(x_0, w')$, and $\eta(r'_0, w')$ are defined. Hence

$$\delta_{A,A'}\big(q_0, (\bar{w}, \bar{w}')\big)! \quad \Leftarrow \quad \eta\big(\eta(r_0, w), \sigma\big)! \ \wedge \ \xi\big(\xi(x_0, w'), \sigma'\big)! \ \wedge \ \eta\big(\eta(r'_0, w'), \sigma'\big)!$$

The right statement is equivalent to $\bar{w}, \bar{w}' \in K$. Thus all the words in $\hat{L}_{A,A'}$ whose length is $n+1$ belong to $L(T_{A,A'})$. This completes the proof. ■

We are now ready to provide a proof of Theorem 15, where the property of $L(T_{A,A'})$ in Lemma 18 plays an important role.

Proof of Theorem 15: ($\Rightarrow$) Suppose that K is not observable under the attack set $\mathcal{A}$. From Proposition 3, there exist $w, w' \in K$, $\sigma \in \Sigma_c$, and $A, A' \in \mathcal{A}$ such that

$$AP(w) \cap A'P(w') \neq \emptyset, \ w\sigma \in K, \text{ and } w'\sigma \in L \setminus K.$$

Then Lemma 18 shows that $(w, w') \in L(T_{A,A'})$. Hence $\delta_{A,A'}\big(q_0, (w, w')\big)$ is defined, and (34) in Lemma 17 holds. Define r, x', r' by

$$r := \eta(r_0, w), \quad x' := \xi(x_0, w'), \quad r' := \eta(r_0, w').$$

Then we see from (35) in Lemma 17 that

$$(r, x', r') = \delta_{A,A'}\big(q_0, (w, w')\big),$$

and hence $(r, x', r') \in \mathrm{Ac}_{A,A'}(Q)$. Furthermore, since $w\sigma \in K$ and $w'\sigma \in L \setminus K$, we have (33).

($\Leftarrow$) Conversely, suppose that $A, A' \in \mathcal{A}$, $(r, x', r') \in \mathrm{Ac}_{A,A'}(Q)$, and $\sigma \in \Sigma_c$ satisfy (33). Since $(r, x', r') \in \mathrm{Ac}_{A,A'}(Q)$, it follows that $(r, x', r') = \delta_{A,A'}\big(q_0, (w, w')\big)$ for some $(w, w') \in L(T_{A,A'})$, and hence Lemma 17 shows that

$$r = \eta(r_0, w), \quad x' = \xi(x_0, w'), \quad r' = \eta(r_0, w').$$

In conjunction with (33), this leads to $w\sigma \in K$ and $w'\sigma \in L \setminus K$. On the other hand, since $(w, w') \in L(T_{A,A'})$, it follows from Lemma 18 that $w, w' \in K$ and $AP(w) \cap A'P(w') \neq \emptyset$. Thus Proposition 3 shows that K is not observable under the attack pair $\{A, A'\}$. This completes the proof. ■

V. Conclusion

We studied supervisory control for DESs in the presence of attacks. We defined a new notion of observability under attacks and proved that this notion combined with conventional controllability is necessary and sufficient for the existence of a supervisor enforcing the specification language despite the attacks. Furthermore, we showed that the usual notion of observability can be used to test the new notion of observability in the case of insertion-removal attacks. The automaton representation and the observability test were extended for replacement-removal attacks. In [35], we also define normality under attacks and discuss the existence of a maximally permissive supervisor. An important direction for future work is to combine robustness against attack with confidentiality and integrity in supervisory control for DESs.

References

[1] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, Shacham, S. H. Savage, K. Kocher, A. Czeskis, F. Roesner, and T. Kohno, "Comprehensive experimental analyses of automotive attack surfaces," in *Proc. USENIX Security Symposium*, 2011.

[2] A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Unmanned aircraft capture and control via GPS spoofing," *J. Field Robot.*, vol. 31, pp. 617–636, 2014.

[3] J. Slay and M. Miller, "Lessons learned from the Maroochy water breach," in *Proc. Critical Infrastructure Protection*, vol. 253, 2007, pp. 73–82.

[4] J. P. Farwell and R. Rohozinski, "Stuxnet and the future of cyber war," *Survival*, vol. 53, pp. 23–40, 2011.

[5] T. De Maiziére, "Die Lage der IT-Sicherheit in Deutschland 2014," Tech. Report, Federal Office for Information Security, 2014. [Online]. Available: http://www.wired.com/wp-content/uploads/2015/01/Lagebericht2014.pdf

[6] P. Falkman, B. Lennartson, and M. Tittus, "Specification of a batch plant using process algebra and petri nets," *Control Eng. Practice*, vol. 17, pp. 1004–1015, 2009.

[7] X. Zhao, P. Shi, and L. Zhang, "Asynchronously switched control of a class of slowly switched linear systems," *Systems Control Lett.*, vol. 61, pp. 1151–1156, 2012.

[8] M. Uzam and G. Gelen, "The real-time supervisory control of an experimental manufacturing system based on a hybrid method," *Control Eng. Practice*, vol. 17, pp. 1174–1189, 2009.

[9] H. Fawzi, P. Tabuada, and S. Diggavi, "Secure estimation and control for cyber-physical systems under adversarial attacks," *IEEE Trans. Automat. Control*, vol. 59, pp. 1454–1467, 2014.

[10] Y. Shoukry and P. Tabuada, "Event-triggered state observers for sparse noise/attacks," *IEEE Trans. Automat. Control*, vol. 61, pp. 2079–2091, 2016.

[11] M. S. Chong, M. Wakaiki, and J. P. Hespanha, "Observability of linear systems under adversarial attacks," in *Proc. ACC'15*, 2015.

[12] F. Lin, "Robust and adaptive supervisory control of discrete event systems," *IEEE Trans. Automat. Control*, vol. 38, pp. 1848–1852, 1993.

[13] S. Takai, "Robust supervisory control of a class of timed discrete event systems under partial observation," *Systems Control Lett.*, vol. 39, pp. 267–273, 2000.

[14] A. Saboori and S. H. Zad, "Robust nonblocking supervisory control of discrete-event systems under partial observation," *Systems Control Lett.*, vol. 55, pp. 839–848, 2006.

[15] A. M. Sánchez and F. J. Montoya, "Safe supervisory control under observability failure," *Discrete Event Dyn. System: Theory Appl.*, vol. 16, pp. 493–525, 2006.

[16] A. Paoli, M. Sartini, and S. Lafortune, "Active fault tolerant control of discrete event systems using online diagnostics," *Automatica*, vol. 47, pp. 639–649, 2011.

[17] S. Shu and F. Lin, "Fault-tolerant control for safety of discrete-event systems," *IEEE Trans. Autom. Sci. Eng.*, vol. 11, pp. 78–89, 2014.

[18] S. Amin, X. Litrico, S. Sastry, and A. M. Bayen, "Cyber security of water SCADA systems–Part I: Analysis and experimentation of stealthy deception attacks," *IEEE Trans. Control Systems Tech.*, vol. 21, pp. 1963–1970, 2013.

[19] A. Teixeira, H. Shames, I. Sandberg, and K. H. Johansson, "A secure control framework for resource-limited adversaries," *Automatica*, vol. 51, pp. 135–148, 2015.

[20] S. Takai and Y. Oka, "A formula for the supremal controllable and opaque sublanguage arising in supervisory control," *SICE J Control Meas. System Integr.*, vol. 1, pp. 307–311, 2008.

[21] J. Dubreil, P. Darondeau, and H. Marchand, "Supervisory control for opacity," *IEEE Trans. Automat. Control*, vol. 55, pp. 1089–1100, 2010.

[22] A. Saboori and C. N. Hadjicostis, "Opacity-enforcing supervisory strategies via state estimator constructions," *IEEE Trans. Automat. Control*, vol. 57, pp. 1155–1165, 2012.

[23] Y.-C. Wu and S. Lafortune, "Synthsis of insertion functions for enforcement of opacity security properties," *Automatica*, vol. 50, pp. 1336–1348, 2014.

[24] D. Thorsley and D. Teneketzis, "Intrusion detection in controlled discrete event systems," in *Proc. 45th IEEE CDC*, 2006.

[25] S.-J. Whittaker, M. Zulkernine, and K. Rudie, "Toward incorporating discrete-event systems in secure software development," in *Proc. ARES'08*, 2008.

[26] N. Hubballi, S. Biswas, S. Roopa, R. Ratti, and S. Nandi, "LAN attack detection using discrete event systems," *ISA Trans.*, vol. 50, pp. 119–130, 2011.

[27] L. K. Carvalho, Y.-C. Wu, R. Kwong, and S. Lafortune, "Detection and prevention of actuator enablement attacks in supervisory control systems," in *Proc. WODES'16*, 2016.

[28] S. Xu and R. Kumar, "Discrete event control under nondeterministic partial observation," in *Proc. IEEE CASE'09*, 2009.

[29] T. Ushio and S. Takai, "Nonblocking supervisory control of discrete event systems modeled by Mealy automata with nondeterministic output functions," *IEEE Trans. Automat. Control*, vol. 61, pp. 799–804, 2016.

[30] R. Cieslak, C. Desclaux, A. S. Fawaz, and P. Varaiya, "Supervisory control of discrete event processes with partial observations," *IEEE Trans. Automat. Control*, vol. 33, pp. 249–260, 1988.

[31] F. Lin and W. M. Wonham, "On observability of discrete-event systems," *Inform. Sci.*, vol. 44, pp. 173–198, 1988.

[32] J. N. Tsitsiklis, "On the control of discrete-event dynamical systems," *Math. Control Signals Systems*, pp. 95–107, 1989.

[33] P. J. Ramadge and W. M. Wonham, "The control of discrete event systems," *Proc. IEEE*, vol. 77, pp. 81–98, 1989.

[34] C. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*. Boston, MA: Kluwer, 1999.

[35] M. Wakaiki, P. Tabuada, and J. P. Hespanha, “Supervisory control of discrete-event systems under attacks,” Tech. Report, University of California, Santa Barbara, 2016. [Online]. Available: http://www.ece.ucsb.edu/~hespanha/published/DES_under_attack_ver8_19_Tech_note.pdf